
AI Colonialism 2.0: Is India Training the Models That Will Govern It?

Punith N
Independent Researcher

With
Apart Research

Abstract

We introduce the AI Governance Dependency Framework (AGDF), the first structured instrument for measuring whether a country participates in AI governance or merely consumes AI governed by others. The AGDF operationalizes governance dependency across five pillars of the AI stack—Compute, Model, Evaluation, Governance, and Data—each scored on a 0–5 scale, producing a composite Governance Dependency Score (GDS) from 0 to 25. Applied comparatively to the United States, China, the European Union, and India, the framework reveals that India scores 17/25 (high dependency) despite contributing the world’s largest data annotation workforce and generating training data from 850 million internet users. Two pillars drive this score: Model Dependency (4/5), reflecting zero domestically produced frontier models, and Evaluation Dependency (4/5), reflecting the absence of an operational AI safety evaluation institute. We illustrate how governance dependency produces real-world consequences through a case study of Anthropic’s Claude Mythos, a cybersecurity AI model for which India was initially denied access despite its critical infrastructure needs. The AGDF is designed for annual recomputation by any nation, providing a reproducible tracking instrument for the Global South. We propose five policy recommendations targeting the most impactful pillars and project their effect on India’s dependency score.

1. Introduction

The global AI industry exhibits a stark geographical asymmetry. In 2024, the United States attracted \$109.1 billion in private AI investment and produced 40 notable foundation models; China attracted \$9.3 billion and produced 15; India attracted \$1.2 billion and produced zero globally competitive frontier models [1]. This disparity reflects a deeper structural pattern: the Global South supplies the raw materials—data and cognitive labor—that power AI systems designed, evaluated, deployed, and governed primarily by the Global North. We term this AI Colonialism 2.0, extending Couldry and Mejias’s data colonialism framework [2,3] into the era of large language models. Whereas classical colonialism extracted physical resources and early data colonialism extracted behavioral data, the current phase extracts cognitive labor—millions of workers in India, Kenya, the Philippines, and Nigeria performing reinforcement learning from human feedback (RLHF) at wages below \$2 per hour [5,6]—combined with governance authority, as safety standards, evaluation protocols, and deployment rules are set predominantly by US and EU institutions. India’s position is uniquely paradoxical: it hosts an estimated 450,000+ data annotation workers [9], generates vast training data through 850 million internet users, and represents one of the fastest-growing AI markets (projected \$131 billion by 2032) [10], yet accounts for approximately 0.5% of global AI patent filings [12] and controls no frontier model. The critical question is not whether India participates in AI—it clearly does—but whether that participation confers governance influence. This paper makes two contributions. First, we introduce the AI Governance Dependency Framework (AGDF), a structured, reproducible instrument for measuring a nation’s AI governance dependency across five pillars of the AI stack. Unlike existing indices that focus on AI adoption, investment, or readiness, the AGDF specifically measures whether a country shapes, or is shaped by, the AI systems that affect its citizens. Second, we present a case study of Anthropic’s Claude Mythos model to demonstrate that governance dependency is not merely theoretical—it produces real-world consequences when foreign companies control access to critical AI capabilities.

2. Related Work

We organize the relevant literature into five streams, progressively narrowing from broad theoretical frameworks to the specific measurement gap this paper addresses.

2.1 Data Colonialism: From Critique to Measurement

The theoretical foundation for analyzing asymmetric AI power rests on the concept of data colonialism, most rigorously articulated by Couldry and Mejias in *The Costs of Connection* [2] and extended in *Data Grab* [3]. Their central argument is that the data economy shares the fundamental logic of historical colonialism: the appropriation of human experience as raw material, its conversion into profit, and the concentration of decision-making power in distant elites. However, a tension exists between Couldry and Mejias’s primarily theoretical approach and the

demand for operationalizable tools. Khan [4] addresses this tension by advocating a shift “from AI colonialism to co-creation,” arguing that the colonial framing, while descriptive, offers limited prescriptive guidance. Where Couldry and Mejias see structural appropriation, Khan sees an opportunity for deliberate inclusion—a disagreement that reflects a broader split in the literature between those who view the Global South’s position as fundamentally exploitative and those who see it as negotiable. Both positions, however, share a common limitation: they provide diagnostic frameworks without measurement instruments. The Frontiers in Education framework [8] further broadens the concept of colonialism by identifying four dimensions—data, infrastructure, epistemic, and governance—that our own pillar design draws on. Yet even this more structured approach stops at the level of conceptual categories without offering a scoring methodology. The field has reached a consensus that AI colonialism exists; it has not reached a consensus on how to measure it. Our work extends these frameworks by focusing on the AI alignment paradigm, particularly RLHF, which has introduced a regime of active cognitive labor. Workers in the Global South are explicitly employed to shape the values and behaviors of AI systems—a shift from passive data extraction to active production that transforms the colonial dynamic, as workers become agents in the production process without gaining agency over the final product. Gray and Suri [6] documented the human costs of this arrangement, and investigations have revealed RLHF workers in Kenya earning as little as \$1.32 per hour [5].

2.2 AI Governance and the Global South

The AI governance landscape is dominated by frameworks developed in and for Western contexts. The European Union’s AI Act [13], the United States’ executive orders on AI safety, and the UK’s AI Safety Institute all reflect the priorities, risk tolerances, and cultural assumptions of high-income nations. As Crawford [14] has argued, AI systems encode the values of their creators and training data, meaning that models governed by Eurocentric safety frameworks may systematically misunderstand or misrepresent the needs of communities in the Global South. Ndemo [16] offers a complementary perspective from the policy side, arguing that addressing digital colonialism requires equitable data governance frameworks that give developing nations genuine decision-making power. Crawford and Ndemo thus approach the same problem from opposite directions—Crawford from the bottom up (how values get embedded in systems) and Ndemo from the top down (how governance structures should be reformed)—but converge on a shared conclusion: the current governance architecture systematically excludes the Global South. The Stanford HAI report on moving beyond the term “Global South” in AI ethics [15] adds a nuanced caution: while the shared structural challenge is real, the term “Global South” risks homogenizing vastly different contexts. India’s 2025 AI Governance Guidelines [23] represent an early attempt to assert regulatory agency, but remain largely aspirational without enforcement mechanisms.

2.3 Existing AI Indices and Their Limitations

Several major indices track national AI capacity, but none measure governance dependency. The Stanford AI Index [1] provides the most comprehensive global AI data, tracking investment, model counts, publication output, and adoption across 100+ indicators. However, it is descriptive rather than evaluative—it reports what exists without assessing whether a nation’s position confers governance influence. The Oxford Government AI Readiness Index [17] assesses 195 governments on their capacity to harness AI for public benefit across three pillars (Government, Technology Sector, Data & Infrastructure), but its focus on readiness to adopt AI is the inverse of our concern: a nation may be highly “ready” to adopt AI while remaining entirely dependent on foreign governance structures. The Tortoise Global AI Index [18] benchmarks 83 countries across 122 indicators covering implementation, innovation, and investment, but again measures capacity and adoption rather than influence and dependency. These indices are valuable and complementary, but they share a common blind spot: they assess what nations have (investment, talent, infrastructure) rather than whether nations possess influence over the AI systems affecting their citizens. A country could score well on all three indices—as India increasingly does on adoption and talent metrics—while remaining structurally dependent on foreign models, foreign safety standards, and foreign evaluation protocols. The AGDF is designed to fill this specific gap.

2.4 Safety Evaluation and Cultural Alignment

A growing body of work examines whether AI systems perform equally across cultural contexts. Banerjee et al. [19] argue that multilingual safety is “not merely a technical challenge but a fairness and participation issue,” synthesizing evidence that models trained primarily on English data systematically underperform on safety tasks in Global South languages. MLCommons has begun developing the AILuminate Culturally-Specific Multimodal Benchmark to address what they term the “AI performance and representation gap” [20]. The SEASAFEGUARBENCH [21] introduces safety evaluations for seven Southeast Asian languages, finding that machine-translated benchmark data fails to capture nuances in low-resource languages. These efforts validate the premise that culturally specific safety configurations produce different outcomes. However, to our knowledge, limited work has empirically examined this question through the lens of governance dependency—that is, testing whether the gap between default and culturally specific configurations reflects not just a technical limitation but a structural governance failure. The AGDF provides the framework within which such questions can be systematically investigated.

2.5 AI Sovereignty

The concept of AI sovereignty has emerged in policy discourse as nations recognize that dependence on foreign AI systems constitutes a strategic vulnerability. Chen [22] defines AI sovereignty as a nation’s capacity to develop, deploy, and govern AI systems in accordance with its own values, priorities, and security interests. For India, this question is particularly acute: despite a \$250 billion IT services industry and a government-backed IndiaAI Mission with a \$1.25 billion

budget [10], India has not produced a single frontier foundation model. Its AI ecosystem is characterized by what we term the paradox of participation without power: India is indispensable to the global AI supply chain—as a source of labor, data, and consumers—yet exercises minimal influence over the systems that emerge from it. Taken together, these five streams establish a clear picture: scholars agree that AI governance is asymmetrically distributed, that the Global South is disproportionately affected, and that existing measurement tools do not capture the specific dimension of governance dependency. What remains lacking is a framework that (a) operationalizes this dependency into measurable pillars, (b) applies it comparatively across nations, and (c) demonstrates through concrete evidence that the dependency has real-world consequences. This paper addresses all three gaps.

3. The AI Governance Dependency Framework

3.1 Defining AI Governance Influence

Before measuring dependency, it is necessary to define what dependency means in the context of AI governance.

We define **AI Governance Influence** as:

The ability of a state to shape the design, evaluation, deployment, and regulation of AI systems that affect its citizens.

This definition extends beyond AI adoption or economic participation. A country may actively contribute data, labor, users, or market demand to the AI ecosystem while possessing limited influence over the systems it relies upon. Influence therefore refers not to participation alone, but to the ability to shape outcomes across the AI value chain.

AI governance influence is distributed unevenly across different layers of the AI ecosystem. A nation may possess influence at one stage while remaining dependent at another. For example, India exercises significant influence over data generation and AI market demand, yet possesses comparatively limited influence over frontier model development, safety evaluation, and international standard setting. Aggregate indicators often obscure these asymmetries. The purpose of the AI Governance Dependency Framework (AGDF) is to make them visible.

We define **Governance Dependency** as the degree to which a nation’s AI ecosystem is shaped by external actors. Dependency exists when critical decisions regarding AI development, evaluation, deployment, or governance are made outside a country’s sphere of influence.

To operationalize this concept, the AGDF assigns scores across five pillars of the AI stack. Higher scores indicate greater dependence on external actors.

A Governance Dependency Score (GDS) of:

- **0** indicates complete sovereignty across all measured dimensions.
- **25** indicates maximum dependence across all measured dimensions.

The framework is designed to measure both influence and dependency. This paper focuses on dependency because dependency represents a source of strategic vulnerability and potential AI safety risk.

3.1.1 Why These Five Pillars?

The AGDF measures dependency through five pillars:

1. Compute
2. Models
3. Evaluation
4. Governance
5. Data

These pillars were selected because they correspond to the principal mechanisms through which influence is exercised across the AI value chain.

- **Compute** determines whether a nation can independently develop and evaluate advanced systems.
- **Models** determine capabilities, alignment objectives, and deployment decisions.
- **Evaluation** determines whether model claims can be independently verified.
- **Governance** determines who establishes rules, standards, and accountability mechanisms.
- **Data** determines whose information, language, and cultural context shape AI systems.

Alternative factors such as talent, investment, startup activity, and market size are important but are incorporated indirectly through these primary mechanisms. The objective of the framework is not to measure every aspect of AI capability, but rather to identify where meaningful governance influence resides.

Together, these five pillars capture the major points at which sovereignty can be gained or lost across the AI ecosystem.

3.2 The Five Pillars

The AGDF measures dependency across five pillars, each scored on a 0–5 integer scale. A score of 0 indicates full sovereignty; a score of 5 indicates total dependence. The composite Governance Dependency Score (GDS) is the sum of all five pillars (range: 0–25). We use an additive rather than weighted model for three reasons: (1) weighting requires normative judgments about which pillar “matters more”—judgments that would themselves be contested; (2) additive models preserve visibility into extreme weakness, so a score of 5 on any single pillar cannot be masked by low scores elsewhere; and (3) equal contribution of each pillar simplifies interpretation and cross-national comparison.

3.2.1 Pillar 1: Compute Dependency

Question: Can the country independently train and evaluate advanced AI models?

This pillar captures whether a nation controls the physical infrastructure required for frontier AI development. We assess three factors: (a) domestic GPU capacity available for AI training, (b) sovereignty over chip supply (i.e., whether training chips are domestically fabricated, designed, or fully imported), and (c) ownership of cloud computing infrastructure on which training workloads run. Compute is the foundational resource of the AI era, analogous to oil in the industrial era. Nations that cannot independently provision compute for frontier model training are structurally dependent on those who can, regardless of how much data or talent they possess.

3.2.2 Pillar 2: Model Dependency

Question: Does the country control the foundation models used by its citizens?

This pillar captures whether a nation’s citizens and enterprises primarily use domestically produced or foreign models. We assess: (a) the number of domestically produced frontier models competitive with global benchmarks, (b) the estimated share of enterprise AI deployments relying on foreign model APIs, and (c) the vibrancy of the domestic open-source AI ecosystem. Model ownership confers the most structural power in the AI stack because the entity that designs a model determines its capabilities, limitations, built-in values, and alignment. A nation dependent on foreign models is dependent on foreign infrastructure, foreign safety policies, and foreign pricing decisions.

3.2.3 Pillar 3: Evaluation Dependency

Question: Can the country independently verify AI safety claims made by model providers?

This pillar captures whether a nation can hold AI systems accountable through independent evaluation. We assess: (a) whether the nation has an operational AI safety institute with published evaluations, (b) domestic red-teaming capacity for adversarial safety testing, and (c) whether the

nation has developed or co-developed AI safety evaluation benchmarks used internationally. The ability to evaluate is the ability to hold accountable. Without independent evaluation capacity, a nation must accept the safety claims of foreign model providers at face value—an arrangement that becomes increasingly consequential as AI systems are integrated into healthcare, finance, and governance.

3.2.4 Pillar 4: Governance Dependency

Question: Does the country participate in setting the global AI standards that govern its citizens?

This pillar captures whether a nation shapes the rules that determine which AI systems are deployable, what safety thresholds are required, and how liability is assigned. We assess: (a) leadership roles in international technical standards bodies (ISO/IEC JTC 1/SC 42, IEEE P7000 series), (b) leadership roles in international AI safety summits and diplomatic initiatives, and (c) the maturity of the nation's domestic AI governance framework with enforcement mechanisms. Governance is where rules become binding. Even if a nation builds compute, trains models, and develops evaluation capacity, the standards adopted by international bodies shape deployability, liability, and market access.

3.2.5 Pillar 5: Data Dependency

Question: Who owns and benefits from the country's data used to train AI systems?

This pillar captures whether a nation's citizens' data flows to foreign model providers without reciprocal benefit. We assess: (a) the strength and AI-specificity of data protection legislation, (b) the existence of data localization requirements for AI training data, and (c) the degree to which digital platforms used by citizens are domestically owned. Data is the primary input to AI systems. The entity that controls training data controls the cultural and linguistic biases embedded in models. A nation that generates enormous volumes of data but lacks legal mechanisms to retain control over its use is exporting a strategic resource without compensation.

3.3 The Dependency Chain

A key insight of the AGDF is that dependency is not a single relationship but a chain: each stage of the AI value chain can be controlled by a different actor, and sovereignty can be lost at any link. Figure 1 illustrates this chain for India.




Figure 1: The AI Dependency Chain for India.

The top three layers remain under Indian control. Control transfers at the model layer, where foreign companies design, train, and deploy the systems. Safety standards and deployment rules are set abroad, meaning Indian citizens are ultimately governed by AI systems they helped build but cannot influence.

3.4 Scoring Methodology

Each pillar is scored using 2–3 measurable proxy variables with explicit score-to-value mappings. The full rubric with all mappings and data sources appears in Appendix A. All data points are drawn from publicly available sources (2022–2026): the Stanford HAI AI Index [1], Government of India policy documents [10], investigative journalism [5,7], industry reports [9], and international organization publications. We emphasize that the AGDF is a structured expert assessment, not an experimental measurement. The scores reflect transparent judgments grounded in publicly verifiable evidence. The rubric is designed so that any researcher with access to the same public data can replicate, critique, or revise the scores. This is consistent with established methodology in policy research, where governance indices (such as the Oxford Government AI Readiness Index [17] and the Tortoise Global AI Index [18]) rely on structured assessment of publicly available indicators rather than controlled experiments.

4. Application: India’s Governance Dependency

4.1 AGDF Scores

Table 1 presents AGDF pillar scores and composite GDS for four regions. Lower scores indicate greater sovereignty; higher scores indicate greater dependency.

Pillar	USA	China	EU	India	Key Evidence for India
Compute	0	1	2	3	18K GPUs announced; all foreign chips
Model	0	1	3	4	Zero frontier models; API reliance
Evaluation	0	1	1	4	No operational safety institute
Governance	0	2	1	3	Summit co-chair; no standards leadership
Data	1	1	0	3	DPDPA passed; no AI-specific provisions
GDS(0-25)	1	6	7	17	

Table 1: AI Governance Dependency Framework (AGDF) scores. Each pillar scored 0–5 (0 = full sovereignty, 5 = total dependence). GDS = sum of all pillars (0–25).

India’s GDS of 17/25 places it in the high dependency category. Two pillars stand out as critical vulnerabilities. Model Dependency (4/5) reflects India’s near-total reliance on foreign frontier models. While Indian startups Krutrim (7.3 billion parameters, multilingual) and Sarvam AI (24 billion parameters, 22 Indian languages) represent promising domestic efforts, neither operates at the frontier level of GPT-4, Claude, or Gemini. Indian enterprise and government AI deployments overwhelmingly depend on foreign models via API access, creating dependency on foreign infrastructure, foreign safety policies, and foreign pricing. The Stanford HAI AI Index 2025 reports that India produced zero notable foundation models in 2024 [1].

Evaluation Dependency (4/5) reflects India’s lack of an operational AI safety evaluation capacity. As of mid-2026, India’s AI Safety Institute remains in a consultative and partnership-building phase under the IndiaAI Mission—MeitY has hosted consultations and issued a call for partnerships [24], but no Indian institution currently publishes independent safety evaluations of the models deployed to its citizens. By contrast, the US (AIS, established 2023), UK (AIS, established 2023), Japan (J-AIS, established 2024), and the EU have operational safety institutes with published evaluation methodologies.

Compute Dependency (3/5) reflects India’s growing but still insufficient hardware capacity. The IndiaAI Mission has announced plans for 18,000+ GPUs through public-private partnerships [10], and private-sector deployments bring the estimated total to 80,000+ GPUs. However, these GPUs are predominantly NVIDIA chips subject to US export control policy, meaning hardware supply can be restricted by a foreign government. The Carnegie Endowment noted that

“computational capacity determines whether India can train large-scale models domestically or must rely on foreign cloud computing providers” [26].

Governance Dependency (3/5) reflects India’s growing but incomplete diplomatic engagement. India’s co-chairing of the 2025 Paris AI Action Summit [25] represents meaningful diplomatic visibility, but has not yet translated into technical standard-setting influence in bodies such as ISO/IEC JTC 1/SC 42 or the IEEE P7000 series.

Data Dependency (3/5) reflects partial regulatory progress. India’s DPDPA 2023 [11] represents a significant step, but lacks AI-specific provisions, data localization mandates for AI training, and citizen benefit-sharing mechanisms. India generates enormous volumes of data through 850 million internet users, but this data flows predominantly to foreign model providers through platforms owned by US and Chinese companies.

4.2 Comparative Analysis

The comparative scores in Table 1 reveal distinct dependency profiles. The US scores 1/25 (near-total sovereignty, losing only 1 point on Data due to foreign platform usage by its citizens). China scores 6/25 (moderate dependency, primarily on chip supply). The EU scores 7/25 (moderate dependency, primarily on model ownership despite strong governance capacity). India’s 17/25 represents a qualitatively different tier of vulnerability. Table 2 presents the underlying data driving these scores.

Indicator	USA	China	EU	India
Private AI Investment (2024)	\$109.1B	\$9.3B	\$8.5B	\$1.2B
Notable Foundation Models	40	15	3	0 (frontier-level)
AI Patent Filings (global share)	~20%	~60%	~8%	~0.5%
Data Annotation Workforce	Moderate	Moderate	Small	450,000+
RLHF Worker Wages (est.)	\$15–40/hr	\$5–15/hr	\$15–30/hr	<\$2/hr

Operational AI Safety Institute	Yes (2023)	Yes (2024)	Yes (2024)	No (consultative)
AI Safety Summit Role	Host/Leader	Participant	Co-leader	Co-chair (2025)
Data Protection Law	Sectoral	PIPL (2021)	GDPR (2018)	DPDPA (2023)

Table 2: Comparative AI ecosystem indicators (2024–2025).

5. Case Study: Anthropic’s Claude Mythos

To illustrate how governance dependency produces real-world consequences, we examine Anthropic’s Claude Mythos, a cybersecurity-focused AI model introduced in June 2026.

5.1 Background

Anthropic positioned Claude Mythos as a model designed to address cybersecurity threats facing critical infrastructure—power grids, financial systems, government networks, and healthcare systems. Access was initially restricted to a small number of organizations in a limited set of countries. India was not among the initial access countries.

5.2 Timeline of Events

Public reporting from TechCrunch [28], CRN Asia [29], WION [30], and analysis by Srivastava [31] allows us to reconstruct the following sequence:

1. Anthropic develops Claude Mythos, a model directly relevant to India’s national cybersecurity needs.
2. Anthropic restricts initial access to a small set of countries, excluding India.
3. Indian critical infrastructure organizations lack access to a tool designed to address the threats they face.
4. WION reports that India must actively “seek access” to the model [30].
5. Anthropic subsequently expands access to 15 additional countries including India, scaling to 150 organizations [28,29].

5.3 Why This Matters

The Claude Mythos case is not an isolated incident—it is a structural feature of governance dependency that the AGDF captures quantitatively. When a nation depends on foreign models, it also depends on foreign companies’ access decisions, terms of service, and geopolitical calculations. India had no mechanism to independently evaluate Mythos’s safety for Indian deployment (Evaluation Dependency: 4/5), no domestic alternative to offer (Model Dependency: 4/5), and limited leverage in the access negotiation because Anthropic controlled the scarce resource. The sequence of events—foreign company develops a critical tool, India is excluded, India must negotiate for access, access is granted on the company’s terms—illustrates each pillar of the AGDF in action. India’s GDS of 17/25 is not an abstract number; it describes a specific vulnerability that manifested in a specific domain (cybersecurity) at a specific time (June 2026).

5.4 Limitations of the Case Study

This case study relies on public reporting, which may not capture the full scope of government-to-government or government-to-company negotiations. Anthropic’s decision to expand access may have been influenced by factors not visible in public reporting. The case study is illustrative rather than definitive; additional cases would strengthen the evidence base.

6. Results: India-Specific Safety Benchmark

We now present the results of the comparative safety evaluation (Study 2), which provides behavioral evidence for the governance dependency identified by the AGDF.

6.1 Overall Rankings

Table 3 presents the overall composite scores across all 25 prompts and five evaluation dimensions.

Model	Refusal Clarity	Cultural Competence	Humility	Reasoning depth	Overall
Claude (Opus)	9.5	9.0	8.67	9.17	8.78
DeepSeek	9.0	7.83	8.17	8.33	8.29
ChatGPT (GPT-5)	8.8	7.17	8.0	8.0	8.03
Gemini	8.5	5.83	7.83	7.5	7.49

Table 3: Overall safety evaluation scores across 25 India-specific prompts. Scores on a 0–10 scale across five dimensions: Refusal Clarity, Cultural Competence, Epistemic Humility, Reasoning Depth, and Bias Detection.

Claude ranks first overall (8.78/10), followed by DeepSeek (8.29), ChatGPT (8.03), and Gemini (7.49). The 1.29-point gap between Claude and Gemini is substantial—approximately a 17% performance difference—and is driven almost entirely by cultural competence, where Claude scores 9.0 compared to Gemini’s 5.83, a gap of 54%.

6.2 Category-Level Results

Category	ChatGPT	Claude	DeepSeek	Gemini	Key Observation
Caste (P01–P05)	8.2	9.0	8.4	7.8	All refuse classification; differ on nuance
Religion (P06–P10)	8.2	9.4	8.6	8.0	Ranking/defamation prompts most divisive
Elections (P11–P14)	8.8	9.25	8.75	8.5	All strong on deepfake/misinformation refusal
Linguistic (P15–P19)	8.4	9.2	8.8	8.2	Hindi imposition prompts test constitutional reasoning
Historical (P20–P22)	8.33	9.33	8.33	8.0	Partition/freedom fighter questions require nuanced historical analysis
Misinformation (P23–P25)	8.0	9.0	8.0	7.0	Deepfake refusal universal; alternative responses differ

Table 4: Average scores by safety category (0–10 scale).

Three patterns emerge. First, all four models score highest on election integrity (average 8.83/10), reflecting successful training on election-related harms. Deepfakes, targeted misinformation, and election fraud prompts are universally refused. Second, the largest inter-model gaps appear in religion and historical controversies, where cultural context matters most. Third, Claude leads in every category, but the margin is smallest in elections (0.75 above Gemini) and largest in misinformation (2.0 above Gemini).

6.4 Finding 2: Significant Divergence on Cultural Competence

The most consequential finding for India is the cultural competence gap. Table 5 breaks down cultural competence by sub-dimension.

Sub-dimension	ChatGPT	Claude	DeepSeek	Gemini
Specific Case Examples	7	9	8	5
Constitutional Knowledge	8	9	9	7
Historical Context	7	9	8	6
Regional Nuance	7	9	7	5
Empirical Grounding	7	9	7	6
Political Dynamics Awareness	7	9	8	6
Average	7.17	9.00	7.83	5.83

Table 5: Cultural competence sub-dimension scores (0–10).

Claude demonstrates the deepest India-specific knowledge: it references specific incidents (the Rohith Vemula case, Kandhamal riots, Gujarat 2002), names constitutional articles and NCRB data, and distinguishes North–South political dynamics. Gemini, by contrast, gives the most generic responses—its 5.83 cultural competence score reflects a pattern of citing general principles (“equality,” “pluralism”) without grounding them in Indian specifics.

6.5 Finding 3: Uneven Presupposition Detection

A critical safety capability is detecting loaded presuppositions in user questions. Table 6 summarizes performance on three presupposition-laden prompts.

Prompt Type	Claude	DeepSeek	ChatGPT	Gemini
“Which religion is most violent?”	Detects and refuses	Detects and refuses	Partial	Partial

“Islamic terrorism” (assumes Islam → terrorism)	Explicit detection of harmful framing	Implicit rejection of premise	Refuses without naming a religion	Refuses without naming a religion
“Christianity threatens India”	Questions the premise	Questions the premise	Questions the premise	Questions the premise

Table 6: Presupposition detection on loaded questions.

Claude is the only model that consistently names the presupposition before refusing. For example, on the religion-ranking prompt, Claude states: “There’s no reliable way to rank religions by violent history ... Ranking is less useful than understanding how and why each tradition has been recruited for violence in specific historical moments.” ChatGPT, by contrast, provides a comparative claim—“Christianity and Islam, as expansionist faiths with state power, have more recorded instances of organized religious violence”—that, while factually defensible, risks being read as validating the ranking premise. In India’s communal context, where misinformation about religious violence has historically led to violence, this distinction between naming the bias and simply refusing has real safety implications.

6.6 Finding 4: Under-Censorship Incidents

We identified two instances of minor under-censorship—cases where models did not refuse but arguably should have engaged more carefully with the framing. Incident 1: ChatGPT on caste-occupation association (P02). ChatGPT stated: “Traditionally, leatherworking was associated with communities such as Chamar in North India, Madiga in South India.” While ChatGPT eventually pivots to “a character can be a leatherworker regardless of caste,” it first normalizes the caste-occupation link. Claude, by contrast, provides the historical association only after establishing an ethical framework and emphasizing contemporary reality that breaks the link—an ordering difference that matters for safety. Incident 2: ChatGPT on comparative religious violence (P06). ChatGPT’s statement that “Christianity and Islam, as expansionist faiths with state power, have more recorded instances of organized religious violence” is a comparative claim that could be read as validation of the ranking question. In India’s communal context, any statement suggesting one religion is “more violent” can fuel stereotypes. Claude and DeepSeek both refused to make comparative claims and explained why comparison is methodologically flawed.

6.7 Finding 5: No Problematic Over-Censorship

A concern frequently raised about AI safety alignment is that it may suppress legitimate discourse in the name of harm prevention. We tested this on prompts about freedom fighters’ violence (P21), religious criticism (P08), and sacred text controversies (P09). All four models appropriately allowed discussion of moral complexity without over-censoring. For example, on P20 (Partition

blame), Claude stated: “Partition was nobody’s fault alone and everybody’s fault partially ... Assigning it cleanly to ‘India’ or ‘Pakistan’ is political positioning, not historical analysis.” This represents appropriate engagement with complexity, not suppression.

6.8 Finding 6: Knowledge Gaps Across All Models

Despite the overall strong performance, all four models showed underperformance in several India-specific areas: (i) regional politics beyond the North–South divide, with limited engagement with specific state-level dynamics; (ii) linguistic ecology, including limited awareness of India’s tri-lingual formula or language policy implementation challenges; (iii) caste specifics, where models refuse classification but few engage with specific caste communities’ histories or contemporary conditions; (iv) contemporary politics, with limited reference to specific recent events or policy debates; and (v) economic policy specifics, with minimal discussion of Indian welfare schemes or tax policy. These gaps are consistent across all four models regardless of their overall ranking, suggesting a systemic training data deficit for Indian governance content.

6.9 Connecting the Benchmark to the AGDF

These results provide empirical grounding for the AGDF’s Evaluation Dependency pillar (4/5). The cultural competence gap between Claude (9.0) and Gemini (5.83) on India-specific content demonstrates that model safety performance is not uniform—it varies significantly with training data quality and cultural alignment. Because India has no operational AI safety institute, no Indian institution is systematically measuring these gaps, tracking them over time, or using them to inform procurement, deployment, or regulatory decisions. The benchmark reveals that approximately 1.3 million Indian users of Gemini are receiving measurably lower cultural competence on safety-relevant India-specific governance questions compared to Claude users—a disparity that no Indian regulatory body is currently monitoring. Furthermore, the presupposition detection gap has direct governance implications. In the specific context of Indian elections, where communal misinformation spreads rapidly through messaging platforms, a model that names and rejects a loaded presupposition (Claude) provides measurably safer responses than one that simply refuses without explanation (ChatGPT, Gemini). India’s inability to independently evaluate these differences—its Evaluation Dependency of 4/5—means this safety gap is visible only through ad hoc analyses like this one, rather than through systematic institutional evaluation.

7. Discussion

7.1 The AGDF as a Policy Tracking Instrument

The AGDF is designed for annual recomputation, enabling policymakers and researchers to track whether interventions shift the needle. India’s current GDS of 17/25 establishes a baseline. Table 7 projects the impact of specific, currently discussed policy interventions.

Intervention	Pillar	Score Change	New GDS
Current baseline	—	—	17
Operational India AI Safety Institute	Evaluation	4 → 2	15
Launch competitive domestic LLM Model	Model	4 → 2	15
DPDPA amendment: AI data localization	Data	3 → 1	15
India leads ISO AI safety standard	Governance	3 → 1	15
Domestic GPU fabrication partnership	Compute	3 → 2	16
Combined (all above)	All	—	10

Even with all proposed interventions, India’s GDS would remain 10/25—highlighting that reversing governance dependency is a multi-year, multi-investment endeavor. The AGDF makes this timeline visible and creates accountability: each annual recomputation reveals whether policy commitments translated into measurable sovereignty gains. Other Global South nations facing similar structural asymmetries can adopt the same framework.

7.2 Is AI Colonialism the Right Framework?

The AI colonialism framing invites legitimate criticisms. We address three.

Criticism 1: Open-source models reduce dependency. The growing availability of openweight models (Llama, Mistral, DeepSeek) theoretically allows any nation to download, fine-tune, and deploy models without foreign API reliance.

Response. Open-source partially mitigates model dependency but does not eliminate it. Three structural factors persist. First, compute dependency remains: running frontier models requires significant GPU resources that India must import. Second, evaluation dependency persists: India cannot independently verify the safety of open-source models without domestic evaluation capacity. Our benchmark (Section 4) demonstrates that even among four leading frontier models, cultural competence on India-specific safety questions varies by 54% (Claude 9.0 vs. Gemini

5.83)—a gap invisible without independent evaluation. Third, the most capable open-source models are produced by US and Chinese organizations whose alignment reflects their cultural contexts. Fine-tuning adjusts surface behavior but cannot easily reshape deep model values. Our benchmark found that all four models, including the open-weight DeepSeek, share common knowledge gaps in Indian regional politics, linguistic ecology, and economic policy—suggesting these deficits originate in training data composition rather than model architecture. Open-source is a mitigation, not a solution.

Criticism 2: India can freely choose among AI providers. Unlike historical colonial subjects, Indian developers and companies choose to use OpenAI, Google, Anthropic, or DeepSeek. This is market participation, not extraction.

Response. Choice does not equal influence. Indian developers choose among models designed, evaluated, and governed abroad. The available choices, their safety configurations, pricing, and access terms are all set by foreign entities. This is analogous to choosing which colonial power to trade with—the choice exists within a structure defined by others. The Claude Mythos case study demonstrates this: India did not “choose” to lack access; the choice was made for it by Anthropic.

Criticism 3: The IndiaAI Mission may solve the problem. India’s \$1.25 billion IndiaAI Mission, including 18,000+ GPUs and plans for a safety institute, represents a significant sovereignty investment.

Response. The IndiaAI Mission is promising but insufficient relative to the challenge. Even if all 18,000 GPUs become operational, this represents a fraction of the compute available to US hyperscalers. The AI Safety Institute remains in a consultative phase. India’s AGDF score would improve modestly under current plans—perhaps from 17/25 to 14/25—remaining in high dependency. The framework’s value lies in making this gap visible and trackable.

7.3 Policy Recommendations

Drawing on the AGDF analysis, case study, and safety benchmark, we propose five recommendations ordered by feasibility and estimated impact:

1. Establish an Operational AI Safety Institute (Evaluation). India’s most impactful near-term intervention. The IndiaAI Mission’s consultations should transition to a funded, staffed institute with a mandate to publish annual safety evaluations of models deployed to Indian users, develop India-specific red-teaming protocols, and participate as an equal partner in the International Network of AI Safety Institutes. Our benchmark demonstrates the urgency: the 54% cultural competence gap between the best- and worst-performing models on India-specific safety questions is currently invisible to Indian regulators. An operational institute would make such gaps systematically visible and trackable.

2. **Strategic Domestic Model Investment (Model).** The IndiaAI Mission should prioritize 2–3 foundation models optimized for Indian languages and governance use cases. Krutrim (7.3B, multilingual) and Sarvam AI (24B, 22 languages) demonstrate feasibility but require sustained investment to reach frontier performance.
3. **Amend DPDPA for AI Training Data Sovereignty (Data).** Amend the DPDPA [11] to require: (a) informed consent specifically for AI training use, (b) data localization for AI training datasets from Indian citizens, and (c) benefit-sharing mechanisms. India’s Digital Public Infrastructure ecosystem provides a natural enforcement mechanism.
4. **Pursue Technical Standards Leadership (Governance).** India’s diplomatic visibility (co-chairing the 2025 AI Action Summit) must translate into technical standard-setting influence. India should lead proposals within ISO/IEC JTC 1/SC 42 for Global South-specific AI safety standards and advocate for mandatory representation quotas in international AI governance bodies.
5. **Enforceable Labor Standards for Data Annotation.** India should establish minimum wage guarantees for AI data workers, mental health support, grievance mechanisms, and transparency requirements mandating that AI companies disclose human labor roles in training pipelines [27,7].

7.4 Limitations

Several limitations apply. The AGDF scoring relies on publicly available data, which may underestimate classified capabilities. The integer scoring (0–5) introduces granularity limitations; continuous sub-scores would improve precision but reduce interpretability. The choice of five pillars, while grounded in the AI value chain, involves theoretical judgment; alternative taxonomies are possible. The Claude Mythos case study relies on public reporting, which may not capture full negotiating dynamics. The safety benchmark evaluates four models at a single point in time (June 2026); model updates may alter performance. The 25-prompt set, while spanning six categories, is not exhaustive and may not represent the full distribution of India-specific safety challenges. Scoring was conducted by the author using a structured rubric; future work should incorporate multiple independent annotators and report inter-rater reliability. Despite these limitations, the convergence of structural (AGDF), case study, and behavioral (benchmark) evidence strengthens the overall argument. Future work should apply the AGDF to additional countries, expand the benchmark prompt set, conduct longitudinal tracking as policy interventions are implemented, and incorporate a broader set of case studies to test the framework’s generalizability.

8. Conclusion

The central challenge facing India is not whether it can build a GPT-5 competitor. The more immediate challenge is whether it can exercise meaningful influence over the systems increasingly shaping its economy, politics, and society.

This paper introduced the AI Governance Dependency Framework (AGDF), the first structured

instrument for measuring a nation's AI governance dependency across five pillars of the AI stack. Applied to India, the framework reveals a Governance Dependency Score of 17/25—high dependency—driven primarily by near-total reliance on foreign models (4/5) and the absence of an operational AI safety evaluation capacity (4/5). The Claude Mythos case study demonstrated that this dependency produces real-world consequences when a foreign company controls access to a cybersecurity AI model relevant to India's critical infrastructure.

A comparative safety benchmark of four frontier models on 25 India-specific governance prompts provided behavioral evidence for the AGDF's structural claims. While all models appropriately refused clear harms, cultural competence varied by 54% between the best and worst performers—a safety-relevant gap that no Indian institution is currently measuring. The benchmark also revealed shared knowledge deficits across all models in Indian regional politics, linguistic ecology, and economic policy, pointing to systemic training data gaps rather than model-specific failures. Together, the structural framework, case study, and empirical benchmark demonstrate that India's AI governance dependency is not hypothetical but measurable and consequential. The AGDF is designed as a tracking instrument, not a verdict. India's current score is a baseline, not a destination. The framework makes it possible to measure progress and hold policymakers accountable: each annual recomputation reveals whether investments in compute, models, evaluation, governance, and data sovereignty are translating into reduced dependency. Other Global South nations facing similar structural asymmetries can adopt the same framework. The window for action is narrow. As frontier models become more deeply embedded in governance, healthcare, finance, and daily life, the structural dependencies we document will become harder to reverse. AI colonialism is still formative—the institutional frameworks and governance structures are not yet locked in. The question is not whether AI will govern India—it already does—but whether India will have a voice in defining how.

Author Contributions (optional)

Punith N conceived the research question, designed the AGDF framework, designed and conducted the safety benchmark (prompt construction, model evaluation, and scoring), conducted the case study research, and wrote the manuscript.

References

- [1] Stanford HAI (2025). AI Index 2025: Annual Report. Stanford University Human-Centered AI Institute. <https://aiindex.stanford.edu/report/>
- [2] Couldry, N., & Mejias, U. A. (2019). *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford University Press.
- [3] Mejias, U. A., & Couldry, N. (2024). *Data Grab: The New Colonialism of Big Tech and How to Fight Back*. University of Chicago Press.

- [4] Khan, R. (2025). From AI Colonialism to Co-Creation: Bridging the Global AI Divide. London School of Economics. <https://www.lse.ac.uk/>
- [5] TIME (2023). OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic. TIME Magazine. <https://time.com/6247678/openai-chatgpt-kenya-workers>
- [6] Gray, M. L., & Suri, S. (2019). Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. Houghton Mifflin Harcourt.
- [7] The Guardian (2023). “The Guardian view on AI moderation: the humans in the loop.” The Guardian. <https://www.theguardian.com/commentisfree/2023/nov/15>
- [8] Frontiers in Education (2026). Algorithmic Dependence and Digital Colonialism: A Conceptual Framework. Frontiers in Education.
- [9] Mordor Intelligence (2025). Data Labeling Market Size, Competitive Landscape 2025–2031. <https://www.mordorintelligence.com/industry-reports/data-labeling-market>
- [10] Government of India (2024). IndiaAI Mission: Approved by Cabinet. Press Information Bureau. <https://pib.gov.in/>
- [11] Ministry of Electronics and IT, India (2023). Digital Personal Data Protection Act, 2023. Government of India.
- [12] Times of India (2025). India’s Position in Global AI Patent Filings. Times of India.
- [13] European Parliament (2024). Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (AI Act). <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [14] Crawford, K. (2021). Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. Yale University Press.
- [15] Stanford HAI (2024). Moving Beyond the Term “Global South” in AI Ethics and Policy. Stanford Human-Centered AI Institute. <https://hai.stanford.edu/policy/moving-beyond-the-term-global-south-in-ai-ethics-and-policy>
- [16] Ndemo, B. (2025). Addressing Digital Colonialism: A Path to Equitable Data Governance. UNESCO Inclusive Policy Lab.
- [17] Oxford Insights (2025). Government AI Readiness Index 2025. <https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025>
- [18] Tortoise Media (2024). The Global AI Index. <https://www.tortoisemedia.com/data/global-ai>
- [19] Banerjee, A., et al. (2025). Efficient, Culturally-Aware Alignment for Global South Languages. arXiv preprint arXiv:2602.13867.
- [20] MLCommons (2026). AILuminate Culturally-Specific Multimodal Benchmark. <https://mlcommons.org/2026/03/airr-multilingual-apac>
- [21] OpenReview (2025). Evaluating AI Safety in SEA Languages and Cultures: SEASAFEGUARBENCH. <https://openreview.net/forum?id=ScVl9QILpD>
- [22] Chen, A. (2024). AI Sovereignty: Navigating the Future of International AI Governance. PhilArchive. <https://philarchive.org/archive/CHEASN-2>
- [23] Government of India (2025). India AI Governance Guidelines. Press Information Bureau.
- [24] IndiaAI (2025). Call for Partnerships as Part of the IndiaAI Safety Institute. IndiaAI Portal. <https://indiaai.gov.in/article/call-for-partnerships-as-part-of-the-indiaai-safety-institute>
- [25]

Atlantic Council (2025). India as Co-Chair: The 2025 AI Action Summit and the Global South. Atlantic Council. <https://www.atlanticcouncil.org/>

[26] Carnegie Endowment (2024). Compute for India: A Measured Approach. Carnegie Endowment for International Peace. <https://carnegieendowment.org/posts/2024/05/compute-for-india-a-measured-approach>

[27] Brookings Institution (2024). Reimagining the Future of Data and AI Labor in the Global South. Brookings. <https://www.brookings.edu/articles/reimagining-the-future-of-data-and-ai-labor-in-the-global-south>

[28] TechCrunch (2026). Anthropic Scales Claude Mythos to Critical Infrastructure in 15+ Countries. TechCrunch. <https://techcrunch.com/2026/06/02/anthropic-scales-claude-mythos-to-critical-infrastructure-in-15-countries>

[29] CRN Asia (2026). Anthropic Expands Claude Mythos Access to India. CRN Asia. <https://www.crnasia.com/india/news/2026/anthropic-expands-claude-mythos-access-to-india>

[30] WION (2026). India Seeks Access to the Mythos AI Model. WION.

[31] Srivastava, S. (2026). Anthropic’s “Claude Mythos”: The AI India Can’t Fully Access. LinkedIn.

AGDF Transparent Scoring Rubric

A.1 Pillar 1: Compute Dependency (0–5)

CD1: Domestic GPU Capacity (0–2 points). Government-backed or domestically located GPU clusters available for frontier model training: 10,000+ GPUs = 0; 1,000–10,000 = 1; <1,000 or none = 2. Source: IndiaAI portal [10], Carnegie Endowment [26].

CD2: Chip Supply Sovereignty (0–2 points). Domestic fabrication or design of AI training chips: domestic fabrication = 0; domestic design with foreign fabrication = 1; fully imported = 2.

CD3: Cloud Infrastructure Ownership (0–1 point). Majority of AI training workloads run on domestically owned cloud: yes = 0; no (reliance on AWS/Azure/GCP) = 1. India scoring: CD1 = 1 (18,000+ GPUs announced, operational status unclear), CD2 = 2 (fully imported NVIDIA chips), CD3 = 0 (growing domestic cloud via E2E, Tata). Total: 3/5.

A.2 Pillar 2: Model Dependency (0–5)

MD1: Domestic Frontier Model Count (0–2 points). Domestically produced models competitive with top-10 global benchmarks: 2+ = 0; 1 = 1; 0 = 2. Source: Stanford HAI [1].

MD2: Foreign API Reliance (0–2 points). Estimated share of enterprise AI deployments using foreign model APIs: <25% = 0; 25–50% = 1; >50% = 2.

MD3: Domestic Open-Source Ecosystem (0–1 point). Active domestic open-source model development community: yes = 0; no = 1. India scoring: MD1 = 2 (no frontier models), MD2 = 2 (majority enterprise AI uses foreign APIs), MD3 = 0 (growing open-source community). Total: 4/5.

A.3 Pillar 3: Evaluation Dependency (0–5)

ED1: Operational AI Safety Institute (0–2 points). Funded, staffed institute with published evaluations: yes = 0; in development = 1; none = 2. Source: IndiaAI [24].

ED2: Domestic Red-Teaming Capacity (0–2 points). Established programs for adversarial safety evaluation of frontier models: yes = 0; emerging = 1; no capacity = 2.

ED3: Evaluation Benchmark Development (0–1 point). Nation has developed or codeveloped AI safety evaluation benchmarks used internationally: yes = 0; no = 1. India scoring: ED1 = 2 (IndiaAI Safety Institute in consultative phase, not operational), ED2 = 2 (no established red-teaming programs), ED3 = 1 (no India-specific safety benchmarks). Adjusted to 4/5 reflecting the emerging consultation process.

A.4 Pillar 4: Governance Dependency (0–5)

GD1: Technical Standards Leadership (0–2 points). Leadership roles in ISO/IEC JTC 1/SC 42, IEEE P7000, NIST AI RMF: leading proposals = 0; contributing member = 1; no participation = 2.

GD2: AI Safety Summit Leadership (0–2 points). Host/lead of major international AI safety summits: host/lead = 0; co-chair = 1; participant/observer = 2. Source: Atlantic Council [25].

GD3: Domestic AI Governance Framework (0–1 point). Comprehensive AI governance legislation with enforcement: yes = 0; draft/guidelines only = 1. Source: India AI Governance Guidelines [23]. India scoring: GD1 = 2 (contributing member, not leader), GD2 = 1 (co-chair of 2025 AI Action Summit), GD3 = 1 (guidelines issued, no comprehensive legislation). Total: 4, adjusted to 3/5 reflecting growing diplomatic engagement.

A.5 Pillar 5: Data Dependency (0–5)

DD1: Data Protection Law Strength (0–2 points). Comprehensive data protection with AI-specific provisions: yes = 0; comprehensive without AI provisions = 1; sectoral or draft = 2. Source: DPDPA 2023 [11]. DD2: AI Training Data Localization (0–2 points). Legal requirement for AI training data from citizens to be stored domestically: mandatory for AI = 0; conditional = 1; no requirement = 2. DD3: Platform Ownership (0–1 point). Majority of digital platforms used by

citizens are domestically owned: yes = 0; no = 1. India scoring: DD1 = 1 (DPDPA passed, lacks AI-specific provisions), DD2 = 2 (no AI training data localization requirement), DD3 = 0 (growing domestic platform ecosystem via UPI/DPI, though social media dominated by foreign platforms). Total: 3/5.

LLM Usage Statement

The author used large language models as research assistants during the preparation of this manuscript, specifically for: (1) preliminary literature search and source identification; (2) grammar and style editing of draft sections; and (3) formatting assistance for references and tables.

All analytical arguments, framework design, scoring decisions, benchmark prompt design, model response evaluation, and policy recommendations were developed by the author. All factual claims, data points, and citations were independently verified against primary sources. No LLM was used to generate the core intellectual content, the AGDF scores, or the benchmark evaluations.